CLAIMS

1. A method for solving a constraint satisfaction problem (CSP) defined by a group of variables and constraints applicable to the variables, the method comprising the steps of:

(a) choosing a first state corresponding to a first set of values of the variables;

(b) selecting a hop distance within a state space of the variables responsively to a random distance selection criterion;

(c) choosing a second state corresponding to a second set of the values of the variables, such that the second state is separated from the first state by the hop distance;

(d) comparing a first cost, determined by applying the constraints to the first set of the values of the variables, to a second cost, determined by applying the constraints to the second set of the values of the variables;

(e) if the second cost is closer than the first cost to meeting a condition indicative that the constraints are satisfied, redefining the first state to correspond to the second set of the values of the variables; and

(f) repeating steps (b) through (e) until the second cost meets the condition, whereby the second set of the values of the variables represents a solution of the CSP.

2. The method according to claim 1, wherein choosing the first state comprises determining that a subset of the values of the variables cannot satisfy at least one of the constraints, and choosing the first state at

random from a range of the values of the variables, while eliminating the subset of the values from the range.

3. The method according to claim 1, wherein choosing the first state comprises determining that a subset of the values of the variables are consistent with at least one of the constraints, choosing the first state at random out of a uniform distribution of the subset of the values.

4. The method according to claim 1, wherein comparing the first cost to the second cost comprises computing respective prices of the constraints based on the values of the variables to which the constraints respectively apply, and summing the respective prices to determine the first and second costs.

5. The method according to claim 4, wherein computing the respective prices comprises associating a respective price function with each constraint, such that the price function is zero when the constraint is satisfied and non-zero when one or more of the values of the variables do not satisfy the constraint, and wherein redefining the first state comprises redefining the first state to correspond to the second set of the values of the variables when the second cost is less than the first cost.

6. The method according to claim 1, wherein each of the variables comprises a respective sequence of bits, and

wherein selecting the hop distance comprises selecting a number of the bits to flip in one or more of the variables so as to transform the first state into the second state.

7. The method according to claim 1, wherein redefining the first state comprises learning a characteristic of the hop distance from the first state to the second state, and wherein selecting the hop distance comprises applying the learned characteristic in selecting the hop distance.

8. The method according to claim 7, wherein applying the learned characteristic comprises selecting the hop distance at random out of a plurality of hop distance records.

9. The method according to claim 1, wherein selecting the hop distance comprises providing a predetermined strategy for choosing the second state, and determining at random whether to use the predetermined strategy or the random distance selection criterion in choosing the second state.

10. The method according to claim 9, wherein each of the variables comprises a respective sequence of bits, and wherein providing the predetermined strategy comprises providing a list of the bits to be flipped so as to cover a selected portion of the state space.

11. The method according to claim 1, wherein the variables are characteristic of inputs to a system under test, and comprising (g) applying the inputs to test the system responsively to the second set of the values of the variables.

12. The method according to claim 11, wherein the system comprises an electronic processor, and wherein applying the inputs comprises determining at least one of a command and an address to be input to the processor responsively to the second set of the values of the variables.

13. The method according to claim 1, wherein the variables are characteristic of control parameters of a mechanical system, and comprising (g) generating a command to control the system responsively to the second set of the values of the variables.

14. The method according to claim 1, wherein the variables are characteristic of features of an image containing visual information, and comprising (g) identifying an object in the image based on the features, responsively to the second set of the values of the variables.

15. The method according to claim 1, wherein the variables are characteristic of a natural language input, and comprising (g) parsing the natural language,

responsively to the second set of the values of the variables, so as to interpret the language.

16. The method according to claim 1, wherein the variables are characteristic of a condition, and comprising (g) determining a diagnosis of the condition responsively to the second set of the values of the variables.

17. The method according to claim 1, wherein the variables are characteristic of resources whose use is to be scheduled, and comprising (g) scheduling the use of the resources responsively to the second set of the values of the variables.

18. Apparatus for solving a constraint satisfaction problem (CSP) defined by a group of variables and constraints applicable to the variables, the apparatus comprising a CSP processor, which is arranged to solve the CSP by the steps of:

(a) choosing a first state corresponding to a first set of values of the variables;

(b) selecting a hop distance within a state space of the variables responsively to a random distance selection criterion;

(c) choosing a second state corresponding to a second set of the values of the variables, such that the second state is separated from the first state by the hop distance;

(d) comparing a first cost, determined by applying the constraints to the first set of the values of the

variables, to a second cost, determined by applying the constraints to the second set of the values of the variables;

(e) if the second cost is closer than the first cost to meeting a condition indicative that the constraints are satisfied, redefining the first state to correspond to the second set of the values of the variables; and

(f) repeating steps (b) through (e) until the second cost meets the condition, whereby the second set of the values of the variables represents a solution of the CSP.

19. The apparatus according to claim 18, wherein the processor is arranged to determine that a subset of the values of the variables cannot satisfy at least one of the constraints, and to choose the first state at random from a range of the values of the variables, while eliminating the subset of the values from the range.

20. The apparatus according to claim 18, wherein the processor is arranged to determine that a subset of the values of the variables are consistent with at least one of the constraints, and to choose the first state at random out of a uniform distribution of the subset of the values.

21. The apparatus according to claim 18, wherein the processor is arranged to compute respective prices of the constraints based on the values of the variables to which the constraints respectively apply, and to sum the respective prices to determine the first and second costs.

22. The apparatus according to claim 21, wherein a respective price function is associated with each constraint, such that the price function is zero when the constraint is satisfied and non-zero when one or more of the values of the variables do not satisfy the constraint, and wherein the processor is arranged to redefine the first state to correspond to the second set of the values of the variables when the second cost is less than the first cost.

23. The apparatus according to claim 18, wherein each of the variables comprises a respective sequence of bits, and wherein the processor is arranged to select the hop distance by selecting a number of the bits to flip in one or more of the variables so as to transform the first state into the second state.

24. The apparatus according to claim 18, wherein the processor is arranged to learn a characteristic of the hop distance from the first state to the second state, and to apply the learned characteristic in selecting the hop distance.

25. The apparatus according to claim 24, wherein the processor is arranged to select the hop distance at random out of a plurality of hop distance records.

26. The apparatus according to claim 18, wherein the processor is further arranged to select the hop distance

based on a predetermined strategy for choosing the second state, and to determine at random whether to use the predetermined strategy or the random distance selection criterion in choosing the second state.

27. The apparatus according to claim 26, wherein each of the variables comprises a respective sequence of bits, and wherein the predetermined strategy provides a list of the bits to be flipped so as to cover a selected portion of the state space.

28. The apparatus according to claim 18, wherein the variables are characteristic of inputs to a system under test, and wherein the processor is arranged to determine, responsively to the second set of the values of the variables, the inputs to be applied in testing the system.

29. The apparatus according to claim 28, wherein the system comprises an electronic processor, and wherein the inputs comprise at least one of a command and an address to be input to the processor.

30. The apparatus according to claim 18, wherein the variables are characteristic of control parameters of a mechanical system, and wherein the processor is arranged to generate, responsively to the second set of the values of the of the variables, a command to control the system.

31. The apparatus according to claim 18, wherein the variables are characteristic of features of an image containing visual information, and wherein the processor is arranged to identify, responsively to the second set of the values of the of the variables, an object in the image based on the features.

32. The apparatus according to claim 18, wherein the variables are characteristic of a natural language input, and wherein the processor is arranged to parse the natural language, responsively to the second set of the values of the of the variables, so as to interpret the language.

33. The apparatus according to claim 18, wherein the variables are characteristic of a condition, and wherein the processor is arranged to determine, responsively to the second set of the values of the of the variables, a diagnosis of the condition.

34. The apparatus according to claim 18, wherein the variables are characteristic of resources whose use is to be scheduled, and wherein the processor is arranged to schedule the use of the resources responsively to the second set of the values of the of the variables.

35. A computer software product for solving a constraint satisfaction problem (CSP) defined by a group of variables and constraints applicable to the variables, the product comprising a computer-readable medium, in

which program instructions are stored, which instructions, when read by a computer, cause the computer to solve the CSP by the steps of:

(a) choosing a first state corresponding to a first set of values of the variables;

(b) selecting a hop distance within a state space of the variables responsively to a random distance selection criterion;

(c) choosing a second state corresponding to a second set of the values of the variables, such that the second state is separated from the first state by the hop distance;

(d) comparing a first cost, determined by applying the constraints to the first set of the values of the variables, to a second cost, determined by applying the constraints to the second set of the values of the variables;

(e) if the second cost is closer than the first cost to meeting a condition indicative that the constraints are satisfied, redefining the first state to correspond to the second set of the values of the variables; and

(f) repeating steps (b) through (e) until the second cost meets the condition, whereby the second set of the values of the variables represents a solution of the CSP.

36. The product according to claim 35, wherein the instructions cause the computer to determine that a subset of the values of the variables cannot satisfy at least one of the constraints, and to choose the first state at random from a range of the values of the variables, while eliminating the subset of the values from the range.

37. The product according to claim 35, wherein the instructions cause the computer to determine that a subset of the values of the variables are consistent with at least one of the constraints, and to choose the first state at random out of a uniform distribution of the subset of the values.

38. The product according to claim 35, wherein the instructions cause the computer to compute respective prices of the constraints based on the values of the variables to which the constraints respectively apply, and to sum the respective prices to determine the first and second costs.

39. The product according to claim 38, wherein a respective price function is associated with each constraint, such that the price function is zero when the constraint is satisfied and non-zero when one or more of the values of the variables do not satisfy the constraint, and wherein the instructions cause the computer to redefine the first state to correspond to the second set of the values of the variables when the second cost is less than the first cost.

40. The product according to claim 35, wherein each of the variables comprises a respective sequence of bits, and wherein the instructions cause the computer to select the hop distance by selecting a number of the bits to

flip in one or more of the variables so as to transform the first state into the second state.

41. The product according to claim 35, wherein the instructions cause the computer to learn a characteristic of the hop distance from the first state to the second state, and to apply the learned characteristic in selecting the hop distance.

42. The product according to claim 41, wherein the instructions cause the computer to select the hop distance at random out of a plurality of hop distance records.

43. The product according to claim 35, wherein the instructions further cause the computer to select the hop distance based on a predetermined strategy for choosing the second state, and to determine at random whether to use the predetermined strategy or the random distance selection criterion in choosing the second state.

44. The product according to claim 43, wherein each of the variables comprises a respective sequence of bits, and wherein the predetermined strategy provides a list of the bits to be flipped so as to cover a selected portion of the state space.

45. The product according to claim 35, wherein the variables are characteristic of inputs to a system under test, and wherein the instructions cause the computer to

determine, responsively to the second set of the values of the variables, the inputs to be applied in testing the system.

46. The product according to claim 45, wherein the system comprises an electronic processor, and wherein the inputs comprise at least one of a command and an address to be input to the processor.

47. The product according to claim 35, wherein the variables are characteristic of control parameters of a mechanical system, and wherein the instructions cause the computer to generate, responsively to the second set of the values of the of the variables, a command to control the system.

48. The product according to claim 35, wherein the variables are characteristic of features of an image containing visual information, and wherein the instructions cause the computer to identify, responsively to the second set of the values of the of the variables, an object in the image based on the features.

49. The product according to claim 35, wherein the variables are characteristic of a natural language input, and wherein the instructions cause the computer to parse the natural language, responsively to the second set of the values of the of the variables, so as to interpret the language.

50. The product according to claim 35, wherein the variables are characteristic of a condition, and wherein the instructions cause the computer to determine, responsively to the second set of the values of the of the variables, a diagnosis of the condition.

51. The product according to claim 35, wherein the variables are characteristic of resources whose use is to be scheduled, and wherein the instructions cause the computer to schedule the use of the resources responsively to the second set of the values of the of the variables.